

MIT Open Access Articles

Obligations to Merely Statistical People

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Hare, Caspar. "Obligations to Merely Statistical People." *The Journal of Philosophy* 109.5/6 (2012): 378-390.

As Published: http://www.pdcnet.org/jphil/content/jphil_2012_0109_41400_0378_0390

Publisher: Journal of Philosophy, Inc.

Persistent URL: <http://hdl.handle.net/1721.1/85649>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Obligations to Merely Statistical People

In some cases the morality of action is an inter-personal affair. I am obliged to do something and there is a person to whom I am obliged to do it. I do wrong and there is a person I wrong. Some routine examples: I do wrong, and wrong you, by doing something bad for you, by feeding you contaminated meat. I do wrong, and wrong you, by failing to do something good for you, by ignoring your S.O.S. I do wrong, and wrong you, by violating your rights, by stealing your stuff. I do wrong, and wrong you, by disrespecting you, by brazenly discounting your opinions.

In other cases this may not be so. I am obliged to do something and yet there is no person to whom I am obliged to do it. I do wrong without wronging any person. Some routine examples: I do wrong by failing to vote in an uncompetitive election. I do wrong by thinking impure thoughts of nobody in particular. I do wrong by destroying an item of transcendent beauty, never seen. (I say 'may not be' because all such examples are controversial. It is tempting to think that the morality of action is by nature inter-personal. Why would a moral consideration give rise to something as weighty as an *obligation* if it did not have to do with someone's interests, someone's rights or someone's dignity? So for any putative example of impersonal obligation there is either a hidden victim – perhaps I very mildly wrong each of my fellow citizens by failing to vote, wrong all people of a certain kind by thinking impure thoughts of people like that, wrong all people who might possibly have born witness to the item of transcendent beauty by destroying it – or no obligation at all.)

And some cases fall curiously in the middle. In these cases a moral obligation may seem, *prima facie*, to be between one person and another, but it proves to be difficult, on reflection, to put a finger on exactly who that other person is. One case that has received a

great deal of attention from philosophers is the ‘Same Number Non-Identity Case’¹ (In brief: I knowingly, for no good reason, bring an unhealthy child into existence when I could have brought a numerically distinct, healthy child into existence a month later. It may seem, *prima facie*, as if I have wronged someone, but who exactly have I wronged? The unhealthy child would not have existed if I had acted more responsibly.) I want to talk about a different class of cases, a class of cases that have received less direct attention from philosophers, here. The cases involve actions that seem, *prima facie*, to bring about great harm or benefit, without greatly harming or benefiting any particular people.

Here’s how the paper will go: In section 2 I will try to give an accurate characterization of some representative cases and of the problem they raise. In sections 3 and 4 I will develop what I take to be the best argument that the lack of great of harm or benefit to particular people in these cases matters. In section 5 I will somewhat tentatively suggest that if it does matter, it does not matter very much. Our moral obligations are not significantly weakened by the absence of a person to whom we are obliged.

2. Framing the Problem: Ought we to be Biased Against Merely Statistical People?

Social scientists, bioethicists and political theorists have observed that many of us, when we make decisions about whom to aid, feel a greater sense of obligation toward *identified* people than to *merely statistical* people.² This attitude is sufficiently robust, widespread, and morally suspect, in their view, to earn the name of a *bias*. They have

¹ Given this name and made famous by Derek Parfit in his “Rights, Interests and Possible People,” in *Moral Problems in Medicine*, Samuel Gorowitz ed., Prentice Hall 1976, 369-75, and his *Reasons and Persons*, Oxford University Press 1983, chapter 16.

² The ‘identified vs. statistical’ terminology goes back at least as far as Thomas Schelling. See his “The Life you Save May be Your Own”, in *Problems in Public Expenditure Analysis*, S.B. Chase ed., The Brookings Institution 1968. An influential, systematic discussion of the psychological phenomena was in Karen Jenni and George Loewenstein, “Explaining the ‘Identifiable Victim Effect’”, *Journal of Risk and Uncertainty* 14, 1997, pp. 235-257, For an overview of the subsequent discussion see Norman Daniels, “Reasonable Disagreement about Identified vs. Statistical Victims”, *Hastings Center Report* 41, no. 4, 2011.

invoked it to explain why, for example, we are relatively strongly motivated to rescue boys drowning in ponds as we pass them by, relatively strongly motivated to send money to benefit small girls who have, in the glare of global media, fallen down wells, but relatively weakly motivated to send money to large charities. The drowning boys and trapped girls are identified. The faceless beneficiaries of the large charities are merely statistical. And they have invoked it to explain why, for example, we are relatively strongly motivated to contribute to programs that aim to cure or manage illness (by distributing antibiotics, for example) but relatively weakly motivated to contribute to programs that aim to prevent illness (by distributing vaccines, for example). The beneficiaries of the former programs are identified. The beneficiaries of the latter programs are merely statistical.

Suppose the social scientists, bioethicists and political theorists are right, that we do have this attitude. Is it correct? Do we indeed have a greater obligation to aid an identified person than to aid a merely statistical person? That is the question I want to address here. But it is broad and dirty. It needs to be rendered slim and clean.

For one thing, talk of 'greater' and 'lesser' obligations is obscure. A less obscure way of interpreting the question: When you have a choice between benefiting an identified and a merely statistical person, ought you, other things being equal, to benefit the identified person?

For another thing, talk of 'identified' and 'merely statistical' people is obscure. What is a 'merely statistical person'? A person with 1.8 children? A person composed of numbers? There are no such people!

It may be that the relevant distinction has to do with what you, the possible benefactor, know. Maybe you know a great deal about the one person and very little about the other person. Or maybe you know a great deal about the way in which you will benefit the one person and very little about the way in which you will benefit the other person. Or

maybe, though you don't know very much about either, you are in a position to know a great deal about the one, but not the other person. Or maybe you know that your doing the one thing will benefit one person, while your doing the other thing has an equivalent *expected benefit*, though it may benefit nobody – maybe, for example, you take there to be a one in a hundred chance that it will benefit one hundred people, a ninety nine in a hundred chance that it will benefit nobody.

In these cases there is something subjectively chancy about the process by which you will come to benefit the 'merely statistical person' if you choose to do that. The cases are all very important and interesting³ but here I want to focus on a different sort of case, in which the process is chancy in a more objective way.

Some vocabulary: Say that a process is *counterfactually open* when, supposing that we initiate it, there is no fact of the matter about what its outcome would have been if we had not initiated it. To be precise: let P be the proposition that we initiate the process, and O_1, \dots, O_k be exclusive propositions concerning the relevantly different outcomes of the process. The counterfactual

(CF0) If it had been that P, then it would have been that O_1 or O_2 or...or O_k .

is true, but none of the counterfactuals

(CF1) If it had been that P, then it would have been that O_1 .

(CF2) If it had been that P, then it would have been that O_2 .

...

(CFk) If it had been that P, then it would have been that O_k .

are true.

³ In case you are interested, I discuss them in my forthcoming book *The Limits of Kindness*.

So, for example, processes governed by indeterministic laws may be counterfactually open. Suppose that, in an optics lab, I have the opportunity to fire a photon through a narrow slit. Suppose I don't do it. Now this counterfactual is true:

(CF3) If I had fired the photon then it would have deflected left or deflected right.

But neither of these are true⁴:

(CF4) If I had fired the photon then it would have deflected left.

(CF5) If I had fired the photon then it would have deflected right.

Why? Well, to put the point in the vocabulary of possible worlds, no world in which I fire the photon and it deflects left is relevantly more similar to the actual world than all worlds in which I fire the photon and it deflects right, and vice-versa. Why? Because the physical laws that govern the actual world are no more or less violated in a world in which I fire the photon and it deflects left than in a world in which I fire the photon and it deflects right.

And, for another example, processes governed by deterministic laws whose outcomes are sensitive to differences in initiating conditions over which we have no control may be counterfactually open. I have a quarter in my pocket. I did not flip it just now. This counterfactual is true:

(CF6) If I had flipped the coin then it would have landed heads or tails.

And maybe, if the laws of nature that govern our world are sufficiently deterministic, some counterfactuals with very specific antecedents, like

⁴ I should note that there is some disagreement about how, precisely, to put this point. On one canonic treatment of counterfactuals offered by David Lewis, (CF4) and (CF5) are false. See section 3.4 of David Lewis, *Counterfactuals*, Blackwell 1973. On another canonic treatment of counterfactuals offered by Bob Stalnaker, (CF4) and (CF5) are neither determinately true nor determinately false. They have indeterminate truth value. See Chapter 7 of Bob Stalnaker, *Inquiry* MIT Press 1984. I side with Lewis, and I will talk accordingly (briefly: Because it seems to me that a counterfactual claim is analogous to a claim concerning a story – a story whose details are fixed by the antecedent of the counterfactual and the nature of the actual world. But if the story does not specify, e.g., whether the photon deflects left or right, it is just false to say that, according to the story, the photon deflects left.) But nothing of importance for present purposes turns on this.

(CF7) If I had flipped the coin while its center of gravity was between 1.48318 and 1.48319 meters from the floor, applying between 2.899 and 2.900 Newtons of force to its upper edge at an angle of ... then it would have landed heads.

(CF8) If I had flipped the coin while its center of gravity was between 1.48320 and 1.4321 meters from the floor, applying... then it would have landed tails.

are true. But neither of these counterfactuals is true:

(CF9) If I had flipped the coin then it would have landed heads.

(CF10) If I had flipped the coin then it would have landed tails.

No world in which I flip the coin and it lands heads is relevantly more similar to the actual world (in which I do not flip it, remember) than all worlds in which I flip it and it lands tails, and vice-versa. The antecedent of counterfactuals (CF9) and (CF10), 'If I had flipped the coin...', is *underspecified*.⁵

When there is no fact of the matter about precisely what would have happened if a process had been initiated there may, nonetheless, be precise counterfactual conditional probabilities. So for example, in the coin case this counterfactual is true:

(CF11) If I had flipped the coin then it might, with probability .5, have landed heads.

Say that a counterfactually open process is *evenly weighted* when, supposing that it is not initiated, for each of the relevantly different outcomes, the counterfactual conditional probability of the process having that outcome, if it had been initiated, is the same. Typical coin-flips are indeed evenly weighted, counterfactually open processes.

Vocabulary settled, here are the cases that I would like to focus on

Case 1: Rescue One from a Threat, or Chancily Rescue Another from a Threat

⁵ I discuss conditional underspecification in more detail in my "Obligation and Regret When There is no Fact of the Matter About What Would have Happened if you had not Done What you Did" *Noûs* 45, no. 1, 2011.

You know all this and nothing (of any relevance to your decision) more: Some people by the names of Agnes, Belinda, Cyril, Damien and Edgar are in danger. If you do nothing then they will all die. If you head north then you can intervene to prevent Agnes from dying. If you head south then you can intervene to prevent one of the others from dying – an evenly weighted, counterfactually open process will determine which one (embellish the story as you like). Those are your only options.

Case 2: Rescue One from a Threat, or Rescue Another from a Chancy Threat

You know all this and nothing (of any relevance to your decision) more: Some people by the names of Archie, Bill, Cynthia, Dora and Edith are in danger. If you do nothing then Archie will die, and one of the others will die – an evenly weighted, counterfactually open process will determine which one (again, embellish the story as you like). If you head north then you can intervene to prevent Archie from dying. If you head south then you can intervene to prevent the counterfactually open process from occurring, and thereby prevent any of Bill, Cynthia, Dora and Edith from dying. Those are your only options.

I choose these cases for a reason. Supposing that, typically, the processes by which we come to benefit particular, distant people through charitable donations are counterfactually open,⁶ Case 1 is in one important way analogous to the case where you have a choice between saving a life by diving into a pond, and ‘saving a distant life’ by

⁶ Are they? We have some reason to think so, in spite of the best efforts of charities to persuade us that the well-being of particular individuals is tightly linked to our donations. Our best dynamical models of weather exhibit extreme sensitivity to initial conditions – small differences in earlier states tend to magnify rapidly in later states. Whether it snows in London in January depends on exactly how things are in Santiago in June. If our world is as these models represent it to be then it may be that, for some pairs of counterfactuals like this:

(CF7) If you had donated \$100 to Oxfam in 2007 then, five years later, distant Belinda would not have died of Typhus.

(CF8) If you had donated \$100 to Oxfam in 2007 then, five years later, distant Belinda would have died of Typhus.

neither is determinately true – because their antecedent ‘If you had donated \$100 to Oxfam in 1986...’ is underspecified. Now, we know that our world is not exactly the way that these models represent it to be. Our world is vastly more complex, and may be not governed by deterministic laws. But we have no reason to think that the extra complexity and nomological determinism renders counterfactuals like (CF7) and (CF8) determinately true or false.

donating to a large charity. If you dive into the pond then there is no fact of the matter about what, precisely, would have happened if you had instead contributed to the charity – it may be that there are many dead potential, beneficiaries of the charity who might, with small probability, have lived if you had donated, but there is no dead, potential beneficiary of the charity who would have lived if you had donated.⁷

And, supposing that, typically, the processes by which we come to benefit particular people by way of vaccination programs are counterfactually open, Case 2 is in one relevant way analogous to the case where you have a choice between saving a life by way of curing an illness by way of distributing antibiotics, and ‘saving a life’ by way of preventing illness by way of distributing vaccines. If you distribute the vaccines then there is no fact of the matter about what, precisely, would have happened if you had instead distributed the antibiotics – it may be that there are many healthy, vaccinated people who might, with some small probability, have died if you had instead distributed the antibiotics, but there is no healthy, vaccinated person who would have died if you had distributed antibiotics.⁸

So our slimmer, cleaner question is this: Is it the case that, in both Cases 1 and 2, you ought to head north (and thereby save the ‘identified person’) rather than head south (and thereby save the ‘merely statistical’ person)?

⁷ I should point out that Case 1 is in other important ways disanalogous to the pond vs charity case. In the former case, if you head north, although there is no fact of the matter about exactly who would have lived if you had headed south, there is a fact of the matter about how many of the affected people would have lived if you had headed south. In the latter case, if you dive into the pond, there is no fact of the matter about how many of the affected people would have lived if you had donated to the charity. Do these further differences matter, morally? This is an interesting question. I am inclined to answer no. I explain why in my forthcoming book, *The Limits of Kindness*.

⁸ Again there are other important disanalogies between the cases. In the antibiotics vs vaccine case, supposing that you distribute the vaccine, there is no fact of the matter about how many people would have lived or died if you had distributed the antibiotic.

3. Appealing to Distributional Equity

One, indirect way to approach the question is to think about whether policies that favor heading north in these cases (or desires to head north in these cases, or stable dispositions to head north in these cases...) are in some way good to have, and then appeal to a bridge principle: you ought to act in line with the policies (or desires, or dispositions...) that are in this way good to have. Another, more direct way to approach the question is to focus on the particular act and its aftermath. Is there something to be said for heading north? I want to pursue this other, more direct approach here.⁹

Norman Daniels has argued¹⁰ that we can vindicate a bias towards identified people by appealing to distributional fairness. In these sorts of cases you are distributing a kind of good, a *chance of living*, and one consideration that bears on what you ought to do is whether you are distributing this good in a fair, equitable way.

In Case 2, here are the chances of living that Archie, Bill, Cynthia Dora and Edith will have if you head north or south ('chances' of the more objective kind, yielded by the counterfactually open process):

	<u>Chances of Living in Case 2</u>				
	<i>Archie</i>	<i>Bill</i>	<i>Cynthia</i>	<i>Dora</i>	<i>Edith</i>
<i>You head north</i>	100%	75%	75%	75%	75%
<i>You head south</i>	0%	100%	100%	100%	100%

If you head north then one person will be certain to live, four likely to live. If you head south then one person will have no chance of living, four will be certain to live. The former distribution of chances is surely more equitable, so, other things being equal, you ought to head north.

⁹ I am skeptical of the bridging principles on which the first approach relies. It may be that there are things to be said for policies that, quite generally, favor people who are, in a loose sense, 'identified' over people who are in a loose sense 'merely statistical'. Maybe the effects upon us of our allowing people to die, in the glare of media spotlight, for the sake of people in its shadows, would be chilling. But this has no bearing on what you ought to do.

¹⁰ In his "Reasonable Disagreement About Identified vs. Statistical Victims."

Daniels' underlying idea needs a lot of motivating. Why isn't *living* the relevant good here, not *chances of living*? If I die then I am not significantly better off for having had a high chance of living. If I live then I am not significantly worse off for having had a low chance of living. And why is it not enough for you to equalize your credence of each of them that they will live, thereby distributing *subjective* chances (relative to your own doxastic state) equitably? Why must you distribute the more objective kind of chances equitably?

In any case the idea will not vindicate a general bias toward 'identified' people.¹¹ Here are the chances that Agnes, Belinda, Cyril, Damien and Edgar will have if you head north or south in Case 1 (again we are talking about chances of the more objective kind, yielded by the counterfactually open process).

	<u>Chances of Living in Case 1</u>				
	<i>Agnes</i>	<i>Belinda</i>	<i>Cyril</i>	<i>Damien</i>	<i>Edgar</i>
<i>You head north</i>	100%	0%	0%	0%	0%
<i>You head south</i>	0%	25%	25%	25%	25%

If you head north then one person will be certain to live, four will have no chance of living. If you head south then one person will have no chance of living, four will have some chance of living. The latter distribution is surely more equitable, so Daniels' reasoning would tell in favor of heading south, saving the 'merely statistical' person.

4. Appealing to Person-Affecting, Anti-Aggregationist Principles

Another, in my view more promising, direct way to approach the question is to look carefully at the effects that your heading north or south in these cases will have on particular people. Will your decision be good, neutral or bad for each of the people it affects?

¹¹ I should emphasize that it is not clear to me that Daniels wanted to vindicate a bias against merely statistical people in all the cases we are considering. So this is no criticism of him.

In Case 1, if you head south then your doing so is very bad for Agnes (she is dead, and she would have lived if you had instead headed north), neutral for three people, the losers of the counterfactually open lottery (they are dead, and they would have died in any case, if you had headed south), and very good for one person, the winner of the counterfactually open lottery (that person is alive, and he/she would have died if you had instead headed north.) If you head north, on the other hand, then your doing so is very good for Agnes (alive, when she would have been dead). That is straightforward. How is it for, e.g., Belinda? That is not so straightforward. Whether it is good or bad for Belinda would seem to depend on whether Belinda would have lived or died if you had headed south. But there is no fact of the matter about whether Belinda would have lived or died if you had headed south! All that can truly be said is that Belinda is dead, and if you had headed south then Belinda might, with probability .25, have lived.

We have a strategic choice to make here. We could say that in a case like this, when you act and there is no fact of the matter about whether Belinda would have been better off if you had acted differently, there is no fact of the matter about whether your action was bad or neutral for her. Or we could say that it depends on the counterfactual conditional probabilities. If the probability that she would have been better off is very high, if this is true

(CF12) If you had headed south then Belinda might, with probability .999, have lived.

then your action was very bad for her. But if it is lower, if this is true

(CF13) If you had headed south then Belinda might, with probability .5, have lived.

then your action was less bad for her. And if it is very low, if this is true

(CF14) If you had headed south then Belinda might, with probability .001, have lived.

Then your action was not very bad at all for her.

I say that, if we ever want to be able to talk about the effects of actions on particular people in realistic cases, we had better go down the second route.¹² This means that in Case 1, by heading north, rather than doing something very bad for one person, you do something quite bad for each of four people (Belinda, Cyril, Damien and Edgar).

Now here's a principle with some appeal:

A Person-Affecting Anti-Aggregationist Principle – Distribute Bad Effects

Given a choice between doing something very good for one person, very bad for one person, and doing something very good for one person, quite bad for each of four people, you ought, other things being equal, to do the latter.

Morally speaking, many small bad effects on many people do not add up to one big bad effect on one person. It is better to hurt many people a little than one person a lot.

It follows that you ought to head north in Case 1. You ought to save Agnes, the 'identified' person.

Case 2 is a little different. In that case heading north will be very good for Archie (alive, when he would have died if you had headed south), neutral for three people, the winners of the counterfactually open lottery (alive, when they would have lived in any case if you had headed south), and very bad for one person, the loser of the counterfactually open lottery (dead, when he/she would have lived if you had headed south). Heading south, on the other hand, will be very bad for Archie (dead, when he would have lived if you had headed north) and merely quite good for each of the four others (each of them is alive and, each of them might, with probability .75, have lived if you had headed north.)

Here's another principle, in the same spirit, with some appeal:

¹² As Alan Hayek has pointed out, even in hum-drum, ordinary cases, the most we can hope for is counterfactual conditional probabilities. See his "Most Counterfactuals are False" ms.

A Person-Affecting Anti-Aggregationist Principle – Concentrate Good Effects

Other things being equal, given a choice between doing something very good for one person and very bad for one person, and doing something quite good for four people and very bad for one person, you ought, other things being equal, to do the former.

Morally speaking, many small good effects on many people do not add up to one big good effect on one person. Equity considerations aside, it is better to benefit one person a lot than many people a little.

It follows that you ought to head north in case 2. You ought to save Archie, the 'identified' person.

5. How Far Does Person-Affecting Anti-Aggregationist Reasoning Take Us?

The person affecting principles, *Distribute Bad Effects* and *Concentrate Good Effects*, tell us that, other things being equal, you ought to save 'identified' rather than 'merely statistical' people in the cases we have looked at, because some considerations to do with the effects of your actions on people matter, in the sense that they have some bearing on the morality of action. In Case 1, for example, it matters that by heading north you will not do anything very bad for anybody.

But the principles do not tell us *how much* those considerations matter. How much do they matter? There certainly is something attractive about the idea that they matter a great deal. A confession of which I am not entirely proud: I gave much less money to charity last year than I could have. I think it likely that, if I had given more money, then some distant-from-me person would have been significantly better off than he or she actually is. But I take consolation in the thought that the processes by which we come to benefit particular, distant people through charitable donations are counterfactually open. There is nobody, sick now, who would have been healthy if I had given more money. There is

nobody, dead now, who would have been alive if I had given more money. My not giving more money was just very, very slightly bad for each of a vast multitude of people. If I believed otherwise, if I believed that there was someone sick or dead who would have been alive and healthy if I had been a little more generous, then I would feel very much less comfortable in my skin. It would not matter that I would have no way of knowing who this person was. The thought would cause me shame.

But it is one thing to be consoled by a thought, another to be properly consoled by a thought, yet another to be properly moved to act by a thought. Is the consideration ‘if I do this then I will not have done anything very bad for anybody’ really so morally significant? I am inclined to think not. One way to gauge its significance is to consider a case in which other things are not equal, in which there is something else at stake, something to be said for going the other way. Consider:

Case 3: Rescue One by Way of a Counterfactually Closed Process, or Two by Way of a Counterfactually Open One?

You know all this and nothing more: A million people are in desperate trouble. You can save one of them by heading north or two of them by heading south. If you head north then the lucky one will be selected by way of a fair, practically unpredictable, counterfactually closed process. (e.g. by way of your computer’s ‘random number generator’ – which is governed by an algorithm whose workings you do not know). If you head south then the lucky two will be selected by way of a fair, practically unpredictable, counterfactually open process. Those are your only options.

In this case there remains something to be said for your heading north. If you head north then you will do something only mildly bad for each of 999,999 people – it will be true of each of them that they are dead, but might, with probability $2/1,000,00$, have survived if you had headed south. If you head south, on the other hand, then almost certainly (with

probability $999,998/1,000,000$) you will do something terrible for one person – it will be true of one person that he/she would have survived if you had headed north. But in this case there are also things to be said for your heading south, including this: If you head south then, for each person, your expectation that that person will survive is $2/1,000,000$. If you head north on the other hand then, for each person, your expectation that that person will survive is $1/1,000,000$. Heading south maximizes, for each person, your expectation that that person will survive.

What should you do? It seems to me that you should head south. I say this because it seems to me that a promising way to approach moral questions is to ask what a benevolent, rational person, someone moved only by wanting the best for each one of us, would do in your position. And a benevolent, rational person would head south. Such a person would not be moved by the consideration: “If I do this then I will not have done anything very bad for anybody.” She cares about people, not about the ways in which her actions affect people.

If this is correct then, although there may be consolation, after the fact, to be found in the thought that the victims of our apathy were ‘merely statistical’, it doesn’t follow that, before the fact, we have significantly less of an obligation to aid them. Other things being equal, our obligation to aid two ‘merely statistical’ people is stronger than our obligation to aid one ‘identified’ one.

6. Wrapping Up

In sum: Philosophers have a way of bemoaning the fact that we tend to be very much less motivated to act helpfully in cases where our help might loosely be described as ‘being of merely statistical benefit’ than in cases where our help might loosely be described as ‘benefiting particular, identified people’. But there is an argument for thinking that our motivations have tuned in to something morally important. Many of the former cases share

an interesting feature. The process that begins with our helping or not, and ends with people being better or worse off, is in some way counterfactually open. Either if we don't help there is no fact of the matter about who would have been better off if we had helped, or if we help there is no fact of the matter about who would have been worse off if we had not helped. So either if we don't help there is nobody for whom our action is very bad, or if we help there is nobody for whom our action is very good. And this feature matters. It is not so bad to fail to help when there is nobody for whom your action is very bad, not so great to help when there is nobody for whom your action is very good.

But let's not get too carried away. It is a further question how much the feature matters. How you answer that question will depend on whether you (as I am inclined to do) look, for moral guidance, to the behavior of someone who cares about people, or to the behavior of someone who cares about the ways in which her actions affect people.